

CHAPTER 17

COLLECTIVE INTELLIGENCE

DAVID H. WOLPERT

17.1 MOTIVATION AND BACKGROUND

17.1.1 Collectives

Many systems of self-interested agents have an associated performance criterion that rates the dynamic behavior of the overall system. This chapter presents an introduction to the science of such systems. Formally, **collectives** are defined as any system having the following two characteristics: First, the system must contain one or more agents each of which we view as trying to maximize an associated *private utility*; second, the system must have an associated *world utility* function that rates the possible behaviors of that overall system [1–5]. In practice, collectives are often very large, distributed, and support little, if any, centralized communication and control, although those characteristics are not part of their formal definition.

A naturally occurring example of a collective is a human economy. One can identify the agents and their private utilities as the human individuals in the economy and the associated personal rewards they are each trying to maximize. One could then identify the world utility as the time average of the gross domestic product. (“World utility” per se is not a construction internal to a human economy, but rather something defined from the outside.) To achieve high world utility it is necessary to avoid having the agents work at cross-purposes lest phenomena like liquidity traps or the Tragedy of the Commons (TOC) occur, in which agents’ individually pursuing their private utilities lowers world utility [6]. The obvious way to avoid such phenomena is by modifying the agents’ utility functions to be “aligned” with the world utility. This can be done via punitive legislation. A real-world example of an attempt to do this was the creation of antitrust regulations designed to prevent monopolistic practices.¹

¹ In conventional economics, imposing governmental regulations is viewed as a change in the dynamical laws of the variables constituting the world (e.g., if you perform the proscribed action *A* you now go to jail, whereas that wasn’t the case before). Here instead it is abstracted to be a direct change in the mapping
(continued...)

We do not insist that the agents in a collective really are “trying” to maximize their private utilities, in some teleological sense. We require only that they can be *viewed* that way. This allows us to circumvent the fraught exercise of formulating a definition of what an arbitrary component of some physical system is “trying to do.” This is illustrated with another naturally occurring example of a collective: a spin glass. One can take the agents in such a glass to be the individual spins. Each spin’s “private utility” is (the negative of) its local Hamiltonian (which is determined only by the states of the other spins to which it is coupled). The “world utility” is (the negative of) the Hamiltonian of the entire system. In this example, both the world utility and all of the private utilities are at (local) maxima at equilibrium. This is what allows us to view the spins *as though* they were agents trying to maximize their private utilities.

In addition to such naturally occurring examples, many current artificial systems can be viewed as collectives. For example, the routers in a terrestrial telecommunications network can be viewed as agents in a collective consisting of the entire network. Real-world routers can reasonably be viewed as “trying to maximize” the quality of service accorded to the traffic that crosses them. Hence they can be taken to be the collective’s agents. World utility in this example can then be set to aggregate quality of service of the entire network.

With the advent of ubiquitous cheap computing in the near future, the number of artificial control systems that are collectives should explode. Two obvious examples here are a user’s constellation of multiple wearable computers, and “computational clouds” of computationally enabled household devices. If such distributed systems are not to be extremely brittle, then absent centralized communication and control, the individual components of the system will need to be both autonomous and adaptive. Almost by definition, this means that those components will be using statistical and machine learning techniques of some sort to modify their behavior to try to meet a goal, i.e., to maximize their private utility.² Moreover, in both of these examples, there is an obvious choice of world utility: the satisfaction level of the user(s) of the system.

Other more prosaic examples of artificial collectives will be dynamic job migration and/or data migration across heterogeneous networks of computers. World utility here will be aggregate satisfaction of the network’s users, suitably quantified. As with our previous examples, with centralized control being impractical, for the system to be robust there will have to be adaptive “agents” in the system that can be viewed as trying to maximize associated private utilities. As examples, the individual agents could be the computers in the network, or even the jobs and/or sets of data themselves. Similarly, as soon as associated computational control devices are

between the state of the world and the agent’s utility value, without any change in underlying dynamical laws. (To continue with the example, in this alternative there is no direct introduction of a new variable having to do with some physical jail—rather your utility function is directly changed so that now if you do *A*, your utility value is smaller than if you do not do *A*.)

²When used for this purpose, such techniques are either explicitly or implicitly related to the field Reinforcement Learning (RL) [7–15].

distributed across such systems, many aspects of the management of supply chain, of electric power grid management, of automobile traffic control and automated control of constellations of deployable autonomous vehicles will constitute collectives.

Another broad class of artificial collectives is essentially every system that will involve copious amounts of nanotechnology where many of the nano-scale components in the system have nontrivial associated computational capabilities. This may include everything from continuously deformable wings to smart paint to nano-scale information storage and retrieval systems.

Finally, consider search algorithms that try to find the value of a high-dimensional variable \vec{z} for which a prespecified function f has a large value. Examples of such algorithms are gradient ascent, simulated annealing, evolutionary algorithms, etc. Say we take the final value of f achieved by such an algorithm to be the “world utility” of the entire system’s dynamic history. Assuming each individual component of \vec{z} evolves with the “goal” of maximizing that final value of $f(\vec{z})$, we can view each such component as an agent, with private utility given by the final value of f . (Note that the private utility of an agent depends on variables not directly under the agent’s control, in general.) In this way any search algorithm can be viewed as a collective. However, conventionally, such algorithms use very “dumb” agents (e.g., semi-random agents rather than reinforcement learning [RL]-based agents). They also don’t consider possible modifications to the underlying system, e.g., to the choice of private utilities, that might result in a better value of final value of f . (The design problem of how best to set private utilities is discussed in the next section.) Constructing search algorithms that use techniques of this nature—intuitively, “agentizing” the individual variables of a search problem by providing them with adaptive intelligence—would provide a search algorithm that is immediately parallelizable. Owing to their use of “smart” variables, such algorithms might also lead to substantially better final values of f than conventional search algorithms.

17.1.2 THE DESIGN OF COLLECTIVES

The “forward problem” in the science of collectives is how the precise configuration of the system—including in particular the private utilities of the agents—affects the ensuing behavior, and therefore affects the value of the world utility. In light of the examples above, however, there is another problem that is at least as rich scientifically, but as a practical matter is of more immediate concern. This is the inverse problem: *How should one initialize/update the private utility functions of the individual agents so that the ensuing behavior of the entire collective achieves large values of the provided world utility?* In particular, since in truly large systems detailed modeling of the system is usually impossible, how can we solve this problem in a way that avoids such modeling? Can we somehow solve it if we leverage only the simple assumption that our agents’ learning algorithms are individually fairly good at what they do?

This design problem is related to work in many other fields, including multi-agent systems (MASs), computational economics, mechanism design, reinforcement

learning, statistical mechanics, computational ecologies, (partially observable) Markov decision processes, and game theory. However none of these fields is both applicable in large problems, and directly addresses the *general* inverse problem, rather than a special instance of it. (See [16] for a detailed discussion of the relationship between these fields, involving hundreds of references.)

For example, the subfield of game-theory known as mechanism design might, at first glance, appear to provide techniques for solving the inverse problem. However, mechanism design is almost exclusively concerned with collectives that are at (a suitable refinement of) Nash equilibrium [17–19]. That means that every agent is assumed to be performing *as well as is theoretically possible*, given the behavior of the rest of the system. In setting private utilities and the like on this basis, mechanism design ignores completely the issue of how to design the system so that each of the agents can achieve a good value of its private utility (given the behavior of the rest of the system). In particular it ignores all statistical issues related to how well the agents can be expected to perform for various candidate private utilities. Such issues become crucial as one moves to large systems, where each agent is confronted implicitly with a very high-dimensional RL task.

There are many other issues that arise in bounded rational situations that are not considered by mechanism design since they do not arise when there is full rationality. For example, it is often the case that by “stabilizing” the sequence of actions of some agent p , the other agents, being in a more predictable environment, are able to perform better. Conversely, such enforced stabilization of its actions will often hurt the performance of agent p . Mechanism design almost completely ignores the associated issues of how best to trade off the performance of one agent against that of other agents, or more generally of how best to trade off the degree of rationality of one agent against that of another agent. (Indeed, mechanism design does not even possess a broadly applicable model-independent measure of “degree of rationality.”)

In addition to these problems, many of the techniques derived in mechanism design cannot be applied in numerous application domains, since those techniques are tailored largely to collectives of human beings. In particular, many of those techniques are tailored to the idiosyncrasy of such collectives that their members have hidden variables whose values they do not want to reveal. This idiosyncrasy is reflected in restrictions on the allowed form of the private utilities and the world utility and communication structures among the agents. Indeed, if there were no such restriction, then given the Nash equilibrium presumption of mechanism design, how best to set the private utilities would be a trivial problem: To have the maximum of world utility be a Nash equilibrium, simply set each such private utility to equal the world utility, in a so-called “team game” or an “exact potential game” [20]. To have the analysis be nontrivial, restrictions such as those that apply to the private utilities of human beings are needed.

Not only are the techniques of mechanism design not relevant to many domains because those domains do not have the restrictions assumed in mechanism design, but in addition there are many issues that loom large in such domains about which mechanism design is mute. For example, in computational domains, where the

agents are computer programs each controlling a set of certain variables, we often have some freedom to change how the set of all variables being controlled is partitioned among the agents, and even change the number of such agents. Needless to say, with its focus on human agents, mechanism design has little advice to provide on such issues of how best to define the agents in the first place.

Perhaps the most striking illustration of the shortcoming of mechanism design is the fact that it does not allow for run-time adaptive redesign. For real-world bounded rational agents, the initial design of the system necessarily makes assumptions that invariably are at least partially at variance with reality. To address this, adaptive techniques are employed (e.g., statistical estimation) on the running system to refine initial assumptions, and then modify the design accordingly. Yet almost all of mechanism design has no room for addressing such macro-learning.

There is other previous work that does consider the inverse problem in its broadest sense, and even has each agent explicitly use RL techniques, so that no formal assumption is made in the associated theory that the system is at Nash equilibrium. Despite this use of RL though, in general, in that work the private utilities are set as in a team game. So again, there is no concern for how well the agents can discern how best to act to maximize their utilities. Unfortunately, as intimated previously (and expounded in the following text), ignoring this issue means that the approach scales extremely poorly to large problems. Intuitively, the difficulty is that each agent will have a hard time discerning the echo of its behavior on its private utility when the system is large if that private utility is the world utility; each agent has a horrible “signal-to-noise” problem in such a situation.³

Intuitively, in designing the private utilities of a collective we want them to be aligned with the world utility, in that modifications an agent might make that would increase its private utility also must increase world utility. Fortunately, the equivalence class of such private utilities extends well beyond team-game utilities. In particular, it extends to include utilities that have far better signal-to-noise properties. By using those utilities one can get far better values of world utility than would otherwise be possible. The mathematical theory for how to generate such alternative private utilities is presented in the next section. The following, last section of this chapter then summarizes many experiments that demonstrate that by using those alternative private utilities one can improve performance by up to orders of magnitude, and that the gain in performance grows as the system gets larger.

17.2 The Mathematics of Designing Collectives

In this chapter, attention is restricted to collectives in which the individual agents are pre-fixed, being the players in multistage noncooperative games, with their moves at any single stage in no a priori way restricted by their moves at other times or by the

³To help see this, consider the example of a collective provided by the human economy. A team game in that example would mean that every human gets U.S. GDP as its reward signal, and tries to discern how best to act to maximize that reward signal. At the risk of understatement, this would provide the individual members of the economy with a difficult reinforcement learning task.

moves of the other players. Some techniques for the design of the private utilities in such games are known as the “COllective INtelligence (COIN)” framework [1]. This section presents some of the mathematics necessary to understand that framework. It should be emphasized, however, that the full mathematics of how to design collectives extends significantly beyond what is needed to address such games.⁴

The restricted version of that full mathematics needed to present the COIN framework starts with an arbitrary vector space Z whose elements ζ give the joint move of all players in the collective in some stage. We wish to search for the ζ that maximizes the provided world utility $G(\zeta)$. In addition to G , we are concerned with private utility functions $\{g_\eta\}$, one such function for each variable/player η . We use the notation $\hat{\eta}$ to refer to all players other than η .

We will need to have a way to standardize utility functions so that the numeric value they assign to a ζ only reflects their ranking of ζ relative to certain other elements of Z . We call such a standardization of some arbitrary utility U for player η the *intelligence* for η at ζ with respect to U . Here we will use intelligences that are equivalent to percentiles:

$$\epsilon_{\eta,U}(\zeta) \equiv \int d\mu_{\zeta_{\hat{\eta}}}(\zeta') \Theta[U(\zeta) - U(\zeta')] \quad (17.1)$$

where the Heaviside function Θ is defined to equal 1 when its argument is greater than or equal to 0, and to equal 0 otherwise, and where the subscript on the (normalized) measure $d\mu$ indicates it is restricted to ζ' sharing the same non- η components as ζ .⁵ Intelligence values are always between 0 and 1.

Our uncertainty concerning the behavior of the system is reflected in a probability distribution over Z . Our ability to control the system consists of setting the value of some characteristic of the collective, e.g., setting the private utility functions of the players. Indicating that value of the *global coordinate* by s , our analysis revolves around the following *central equation* for $P(G | s)$, which follows from Bayes' theorem:

$$P(G | s) = \int d\vec{\epsilon}_G P(G | \vec{\epsilon}_G, s) \int d\vec{\epsilon}_g P(\vec{\epsilon}_G | \vec{\epsilon}_g, s) P(\vec{\epsilon}_g | s) \quad (17.2)$$

where $\vec{\epsilon}_g \equiv (\epsilon_{\eta_1, g_{\eta_1}}(\zeta), \epsilon_{\eta_2, g_{\eta_2}}(\zeta), \dots)$ is the vector of the intelligences of the players with respect to their associated private utility functions, and $\vec{\epsilon}_G \equiv (\epsilon_{\eta_1, G}(\zeta), \epsilon_{\eta_2, G}(\zeta), \dots)$ is the vector of the intelligences of the players with respect to G .

⁴That framework encompasses, for example, arbitrary dynamic redefinitions of the “players” (i.e., dynamic reassignments of how the various subsets of the variables comprising the collective across all space and time are assigned to players), as well as modification of the players' information sets (i.e., modification of interplayer communication). See [21].

⁵The measure must reflect the type of system at hand, e.g., whether Z is countable or not, and if not, what coordinate system is being used. Other than that, any convenient choice of measure may be used and the theorems will still hold.

Note that $\epsilon_{\eta, g_\eta}(\zeta) = 1$ means that player η is fully rational at ζ , in that its move maximizes the value of its utility, given the moves of the players. In other words, a point ζ where $\epsilon_{\eta, g_\eta}(\zeta) = 1$ for all players η is one that meets the definition of a game-theory Nash equilibrium.⁶ On the other hand, a ζ at which all components of $\vec{\epsilon}_G = 1$ is a local maximum of G (or more precisely, a critical point of the $G(\zeta)$ surface). So if we can get these two vectors to be identical, then if the agents do well enough at maximizing their private utilities we are assured we will be near a local maximum of G .

To formalize this, consider our decomposition of $P(G | s)$. If we can choose s so that the third conditional probability in the integrand is peaked around vectors $\vec{\epsilon}_g$ all of whose components are close to 1, then we have likely induced large (private utility function) intelligences. If we can also have the second term be peaked about $\vec{\epsilon}_G$ equal to $\vec{\epsilon}_g$, then $\vec{\epsilon}_G$ will also be large. Finally, if the first term in the integrand is peaked about high G when $\vec{\epsilon}_G$ is large, then our choice of s will likely result in high G , as desired.

Intuitively, the requirement that private utility functions have high signal-to-noise arises in the third term. It is in the second term that the requirement that the private utility functions be aligned with G arises. In this chapter we concentrate on these two terms, and show how to simultaneously set them to have the desired form.⁷

Details of the stochastic environment in which the collective operates, together with details of the learning algorithms of the players, are all bundled into the distribution $P(\zeta)$ which underlies the distributions appearing in Eq. 17.2. Note though that *independent of these considerations*, our desired form for the second term in Eq. 17.2 is assured if we have chosen private utilities such that $\vec{\epsilon}_g$ equals $\vec{\epsilon}_G$ exactly for all ζ . Such a system is said to be *factored*. In game-theory parlance, the Nash equilibria of a factored collective are local maxima of G . In addition to this desirable equilibrium behavior, factored collectives also automatically provide appropriate off-equilibrium incentives to the players (an issue rarely considered in the game theory/mechanism design literature).

As a trivial example, any team game in which all the private utility functions equal G is factored [20,24]. However, team games often have very poor forms for term 3 in Eq. 17.2, forms that get progressively worse as the size of the collective grows. This is because for such private utility functions each player η will usually confront a very poor signal-to-noise ratio in trying to discern how its actions affect

⁶ See [12]. Note that consideration of points ζ at which *not* all intelligences equal 1 provides the basis for a model-independent formalization of bounded rationality game theory. This formalization contains variants of many of the theorems of conventional full-rationality game theory. See [22].

⁷ Search algorithms that do not involve game theory (e.g., simulated annealing) can be viewed as addressing how to have term 1 possess the desired form. They do this by trying to ensure that the particular local maximum they find of the function they are searching has a high value of that function. This is the essence of why such algorithms trade off exploration and exploitation. One can combine such term-1-based techniques with the techniques presented in this paper. Intuitively, this amounts to “wrapping” a system using the private utilities derived below in an outer loop that trades off exploration and exploitation. The resultant hybrid algorithm, addressing all three terms, outperforms simulated annealing by over two orders of magnitude [23].




Figure 17.1 Intelligence of agent η at state ζ for utility U is the actual joint move at hand. The x-axis shows agent η 's alternative possible moves (all states ζ' having ζ 's values for the moves of all players other than η). The thick sections of the x-axis show the alternative moves that η could have made that would have given η a worse value of the utility U . The fraction of the full set of η 's possible moves that lies in those thick sections (which is 0.6 in this example) is the intelligence of agent η at ζ for utility U , denoted by $\epsilon_{\eta,U}(\zeta)$.

its utility $g_\eta = G$, since so many other player's actions also affect G and therefore dilute η 's effect on its own private utility function.

We now focus on algorithms based on private utility functions $\{g_\eta\}$ that optimize the signal/noise ratio reflected in the third term, subject to the requirement that the system be factored. To understand how these algorithms work, say we are given an arbitrary function $f(\zeta_\eta)$ over player η 's moves, two such moves ζ_η^1 and ζ_η^2 , a utility U , a value s of the global coordinate, and a move by all players other than η , $\zeta_{\hat{\eta}}$. Define the associated *learnability* by

$$\Lambda_f(U; \zeta_{\hat{\eta}}, s, \zeta_\eta^1, \zeta_\eta^2) \equiv \sqrt{\frac{[E(U; \zeta_{\hat{\eta}}, \zeta_\eta^1) - E(U; \zeta_{\hat{\eta}}, \zeta_\eta^2)]^2}{\int d\zeta_\eta [f(\zeta_\eta) \text{Var}(U; \zeta_{\hat{\eta}}, \zeta_\eta)]}} \quad (17.3)$$

The expectation values in the numerator are formed by averaging over the training set of the learning algorithm used by agent η , n_η . Those two averages are evaluated according to the two distributions $P(U|n_\eta)P(n_\eta|\zeta_{\hat{\eta}}, \zeta_\eta^1)$ and $P(U|n_\eta)P(n_\eta|\zeta_{\hat{\eta}}, \zeta_\eta^2)$, respectively. (That is the meaning of the semicolon notation.) Similarly the variance being averaged in the denominator is over n_η according to the distribution $P(U|n_\eta)P(n_\eta|\zeta_{\hat{\eta}}, \zeta_\eta)$.

The denominator in Eq. 17.3 reflects how sensitive $U(\zeta)$ is to changing $\zeta_{\hat{\eta}}$. In contrast, the numerator reflects how sensitive $U(\zeta)$ is to changing ζ_η . So the greater the learnability of a private utility function g_η , the more $g_\eta(\zeta)$ depends only on the move of player η , i.e., the better the associated signal-to-noise ratio for η . Intuitively then, so long as it does not come at the expense of decreasing the signal, increasing the signal-to-noise ratio specified in the learnability will make it easier for η to achieve a large value of its intelligence. This can be established formally: If scaled,

appropriately g'_η will result in better expected intelligence for agent η than will g_η whenever $\Lambda_f(g'_\eta; \zeta_\eta, s, \zeta_\eta^1, \zeta_\eta^2) > \Lambda_f(g_\eta; \zeta_\eta, s, \zeta_\eta^1, \zeta_\eta^2)$ for all pairs of moves $\zeta_\eta^1, \zeta_\eta^2$ [21].⁸

It is possible to solve for the set of all private utilities that are factored with respect to a particular world utility. Unfortunately, in general it is not possible for a collective both to be factored and to have infinite learnability for all of its players. However, consider *difference* utilities, which are of the form

$$U(\zeta) = \beta[G(\zeta) - \Gamma(\zeta_\eta)] \quad (17.4)$$

Any difference utility is factored [21]. In addition, under usually benign approximations, $\Lambda_f(U; \zeta_\eta, s, \zeta_\eta^1, \zeta_\eta^2)$ is maximized over the of difference utilities for all pairs $\zeta_\eta^1, \zeta_\eta^2$ by choosing

$$\Gamma(\zeta_\eta) = E_f(G(\zeta) \mid \zeta_\eta, s) \quad (17.5)$$

up to an overall additive constant, where the expectation value is over ζ_η . We call the resultant difference utility the *Aristocrat* utility (AU), loosely reflecting the fact that it measures the difference between a player's actual action and the average action. If each player η uses an appropriately rescaled version of the associated AU as its private utility function, then we have ensured good form for both terms 2 and 3 in Eq. 17.2.

Using AU in practice is sometimes difficult, due to the need to evaluate the expectation value. Fortunately there are other utility functions that, while being easier to evaluate than AU, still are both factored and possess superior learnability to the team game utility, $g_\eta = G$. One such private utility function is the *Wonderful Life* Utility (WLU). The WLU for player η is parameterized by a prefixed *clamping parameter* CL_η chosen from among η 's possible moves:

$$WLU_\eta \equiv G(\zeta) - G(\zeta_\eta, CL_\eta) \quad (17.6)$$

WLU is factored no matter what the choice of clamping parameter. Furthermore, while not matching the high learnability of AU, WLU usually has far better learnability than does a team game, and therefore (when scaled appropriately) results in better expected intelligence [5,2,16,25].

Figure 17.2 provides an example of clamping. As in that example, in many circumstances there is a particular choice of clamping parameter for player η that is a “null” move for that player, equivalent to removing that player from the system.

⁸ In many RL algorithms, changing the scale of the utility is exactly equivalent to changing a temperature parameter of the algorithm. Such temperatures have to usually be set via a search process. The result presented here establishes that so long as g'_η has higher learnability than does g_η , the expected intelligence of g'_η at the associated optimal temperature will be higher than that of g_η at its optimal temperature.

$$\begin{array}{c} \eta_1 \\ \eta_2 \\ \eta_3 \\ \eta_4 \end{array} \begin{array}{c} \zeta \\ \left[\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{array} \right] \end{array} \xRightarrow{\text{Clamp } \eta_2 \text{ to "null" }} \begin{array}{c} (\zeta_{\hat{\eta}_2}, \vec{0}) \\ \left[\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{array} \right] \end{array} \xRightarrow{\text{Clamp } \eta_2 \text{ to "average" }} \begin{array}{c} (\zeta_{\hat{\eta}_2}, \vec{a}) \\ \left[\begin{array}{ccc} 1 & 0 & 0 \\ .33 & .33 & .33 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{array} \right] \end{array}$$

Figure 17.2 This example shows the impact of the clamping operation on the joint state of a four-player system where each player has three possible moves, each such move represented by a three-dimensional unary vector. The first matrix represents the joint state of the system ζ where player 1 has selected action 1, player 2 has selected action 3, player 3 has selected action 1 and player 4 has selected move 2. The second matrix displays the effect of clamping player 2's action to the "null" vector (i.e., replacing ζ_{η_2} with $\vec{0}$). The third matrix shows the effect of instead clamping player 2's move to the "average" action vector $\vec{a} = \{.33, .33, .33\}$, which amounts to replacing that player's move with the illegal move of fractionally taking each possible move ($\zeta_{\eta_2} = \vec{a}$).

(Hence the name of this private utility function—cf. the Frank Capra movie.) For such a clamping parameter, assigning the associated WLU to η as its private utility function is closely related to the economics technique of "endogenizing a player's externalities," for example with the Groves mechanism [17,19,26,9].

However it is usually the case that using WLU with a clamping parameter that is as close as possible to the expected move defining AU results in far higher learnability than does clamping to the null move. Such a WLU is roughly akin to a mean-field approximation to AU.⁹ For example, in Fig. 17.2, if the probabilities of player 2 making each of its possible moves were 1/3, then one would expect that a clamping parameter of \vec{a} would be close to optimal. Accordingly, inpractice use of such an alternative WLU derived as a mean-field approximation to AU almost always results in far better values of G than does the "endogenizing" WLU.

Intuitively, collectives having factored and highly learnable private utilities like AU and WLU can be viewed as akin to well-run human companies. G is the bottom line of the company, the players η are identified with the employees of that company, and the associated g_η given by the employees' performance-based compensation packages. For example, for a factored company, each employee's compensation package contains incentives designed such that the better the bottom line of the corporation, the greater the employee's compensation. As an example, the CEO of a company wishing to have the private utilities of the employees be factored with G may give stock options to the employees. The net effect of this action is to ensure that what is good for the employee is also good for the company. In addition, if the compensation packages are highly learnable, the employees will have a relatively easy time discerning the relationship between their behavior and their

⁹ Formally, our approximation is exact only if the expected value of G equals G evaluated at the expected joint move (both expectations being conditioned on given moves by all players other than η). In general though, for relatively smooth G , we would expect such a mean-field approximation to AU, to give good results, even if the approximation does not hold exactly.

compensation. In such a case the employees will both have the incentive to help the company and be able to determine how best to do so. Note that in practice, providing stock options is usually more effective in small companies than in large ones. This makes perfect sense in terms of the formalism summarized above, since such options generally have higher learnability in small companies than they do in large companies, in which each employee has a hard time seeing how his/her moves affect the company's stock price.

17.3 TESTS OF THE MATHEMATICS

As a test of the preceding mathematics, in some of our previous work we used the WLU for distributed control of network packet routing [2]. Conventional approaches to packet routing have each router run a Shortest Path Algorithm (SPA), i.e., each router routes its packets in the way that it expects will get those packets to their destinations most quickly. Unlike with a WLU-based collective, with SPA-based routing the routers have no concern for the possible deleterious side-effects of their routing decisions on the global goal (e.g., they have no concern for whether they induce bottlenecks). We ran simulations that demonstrated that a WLU-based collective has substantially better throughputs than does the best possible SPA-based system [2], even though that SPA-based system has information denied the agents in the WLU-based collective.

In related work, we have shown that use of the WLU automatically avoids the infamous Braess' paradox, in which adding new links can actually decrease throughput—a situation that readily ensnares SPAs [5,27].

In yet other work we have applied the WLU to the problem of controlling communication across a constellation of satellites so as to minimize the importance-weighted loss of scientific data flowing across that constellation.[28] Due to the novelty of this problem domain, we first had to design a baseline distributed control algorithm, one that involves no learning. To minimize the number of confounding distinctions between that baseline algorithm and the collective-based algorithm we investigated, we had that collective “run on top” of the baseline algorithm. The action of each agent in the collective was the determination of fictitious “ghost traffic” that is presented to the baseline algorithm, thereby (hopefully) inducing that baseline algorithm to achieve an even better value of the world utility. (Note that this idea can be applied with most any baseline algorithm and most any distributed RL algorithm.) Again, we achieved a significant increase in performance, in this case relative to the baseline algorithm.

We have also successfully applied the COIN techniques to problems that are explicitly cast as search. These include setting the states of the spins in a spin glass to minimize energy; the conventional bin-packing problem of computer science, and a model of human agents connected in a small-world network who have to synchronize their purchase decisions. We have also successfully applied COIN techniques to the problem of coordinating a set of autonomous rovers so as to maximize the importance-weighted value of a set of locations they visit [29].

Finally, it is worth going into some detail our investigations on variants of congestion games [1,3,25], in particular of a more challenging variant of Arthur's El Farol bar attendance problem [30], sometimes also known as the minority game [31]. In this problem the individual processes making up the collective are explicitly viewed as 'players' involved in a noncooperative game. Each player has to determine which night in the week to attend a bar. The problem is set up so that if either too few people attend (boring evening) or too many people attend (crowded evening), the total enjoyment of the attending players drops. Our goal is to design the private utility functions of the players so that the total enjoyment across all nights is maximized. In this previous work we showed that use of the WLU can result in performance *orders of magnitude* superior to that of team game utilities.

Arthur's bar problem [30] can be viewed as a problem in designing collectives. Loosely speaking, in this problem at each time step each player η decides whether to attend a bar by predicting, based on its previous experience, whether the bar will be too crowded to be "rewarding" at that time, as quantified by a utility function G . The selfish nature of the players frustrates the global goal of maximizing G . This is because if most players think the attendance will be low (and therefore choose to attend), the attendance will actually be high, and vice-versa.

The variants of the bar problem we investigated were all of the following type: There are N players, each picking one out of seven moves every week. Each variant of the game is parameterized by $\ell \in \{1, 2, 3, 4, 5, 6\}$. In a given variant, each move of an agent corresponds to attending the bar on some particular subset of ℓ out of the seven nights of the current week (i.e., given ℓ , each possible move is an 'attendance profile' vertex of the 7-dimensional unit hypercube having ℓ 1's). In each week every player chooses a move. Then the associated private utility values for each player are communicated to that player, and the process is repeated. For simplicity, for each ℓ we chose the seven possible attendance profiles so that if the moves are selected randomly uniformly, the expected resultant attendance profile across all seven nights is also uniform. (For example, for $\ell = 2$, those profiles are (1, 1, 0, 0, 0, 0, 0), (0, 1, 1, 0, 0, 0, 0), etc.)

More formally, the world utility in any particular week is:

$$G(\zeta) \equiv \sum_{k=1}^7 \phi(x_k(\zeta)) \quad (17.7)$$

where $x_k(\zeta)$ is the total attendance on night k ; ζ_η is η 's move in that week; $\phi(y) \equiv y \exp(-y/c)$; and c is a real-valued parameter. Our choice of $\phi(\cdot)$ means that when either too few or too many players attend some night, in some week, world utility G is low.

Since we wished to concentrate on the effects of the utilities rather than on the RL algorithms that use them, we used (very) simple RL algorithms.¹⁰ We would

¹⁰ On the other hand to use algorithms so patently deficient that they have never even been considered in the RL community—like the algorithms used in most of the bar problem literature—would seriously interfere with our ability to interpret our experiments.

expect that even marginally more sophisticated RL algorithms would give better performance. In our algorithm each player η had a 7-dimensional vector giving its estimates of the utility it would receive for taking each possible move. At the beginning of each week, each η picked the night to attend randomly, using a Boltzmann distribution over the seven components of η 's estimated utilities vector. For simplicity, temperature did not decay in time. However, to reflect the fact that each player operated in a nonstationary environment, utility estimates were formed using exponentially aged data: in any week t , the estimate η makes for the utility for attending night i was a weighted average of all the utilities it has received previously when it attended that night, with the weights given by an exponential function of how long ago each such utility was. To form the players' initial training set, we had an initial period in which all moves by all players were chosen uniformly randomly, with no learning.

In these experiments we found once again that use of highly learnable factored private utilities resulted in vastly better performance than use of team game private utilities. Also as usual, we found that the gain in performance grew as the problem grew, reaching orders of magnitude once the system grew to consist of thousands of agents. We always found that AU performed at least as well as WLU with clamping to 0, which is essentially identical to the economics technique of endogenizing externalities. In addition though, for some choices of ℓ , we found that AU performed substantially better than did this alternative, as would be expected based on the formalism presented above.

Finally, the central equation makes numerous other predictions that preliminary experiments seem to bear out. Some of these predictions concern ways to modify the behavior of the collective to try to optimize term 1 as well as terms 2 and 3. (The work in [28] can be viewed as an initial investigation of this issue.) Other predictions are on how to modify a factored private utility so that it is *not* perfectly factored any more, but has undergone such a large gain in learnability that (as quantified in the central equation) overall performance improves [31]. It is worth emphasizing that such beneficial modifications to private utilities are prohibited by the starting premise of the field of mechanism design, that the private utilities must exhibit incentive compatibility.

17.4 CONCLUSION

A collective is any multi-agent system in which each agent adaptively tries to maximize its own private utility, while at the same time there is an overall world utility rating the behavior of the entire system. Collectives are quite common in the natural world, the canonical example being any human organization. In addition, as computing becomes ubiquitous in artificial systems, the number of such systems that constitute collectives will explode.

Associated with any collective is a design problem, of how to configure the system—and in particular how to set the private utilities of the agents—to optimize the world utility. This paper cursorily synthesizes some of the mathematical theory

underpinning this design problem. That theory has now been tested in many different experiments. As summarized above, those tests of the theory have clearly validated it, often resulting in performance up to orders of magnitude superior to traditional techniques from the fields of multi-agent systems and economics/mechanism design. Intuitively, that superiority lies in the fact that these alternative approaches completely ignore the issue of how an agent's ability to maximize a candidate private utility will vary with changes in that private utility. This issue is especially crucial in large systems, in which each agent will face an extremely difficult "signal-to-noise" term in discerning the effects of its actions on its utility unless that utility is crafted carefully.

REFERENCES

1. D. H. Wolpert and K. Tumer. "Optimal Payoff Functions for Members of Collectives." *Advances Complex Systems*, vol. 4, no. 2/3, 265–279, 2001.
2. D. H. Wolpert, K. Tumer, and J. Frank. "Using Collective Intelligence to Route Internet Traffic." In *Advances in Neural Information Processing Systems—11*, pages 952–958. MIT Press, 1999.
3. D. H. Wolpert, K. Wheeler, and K. Tumer. "General Principles of Learning-Based Multi-Agent Systems." In *Proc. Third International Conference of Autonomous Agents*, pages 77–83, 1999.
- 4.
5. K. Tumer and D. H. Wolpert. "Collective Intelligence and Braess' Paradox." In *Proceedings of the 17th National Conference on Artificial Intelligence*, pages 104–109, Austin, TX, 2000.
6. G. Hardin. "The Tragedy of the Commons." *Science*, vol. 162, 1243–1248, 1968.
7. L. P. Kaelbling, M. L. Littman, and A. W. Moore. "Reinforcement Learning: A Survey." *J. Artificial Intelligence Res.*, vol. 4, 237–285, 1996.
8. R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
9. R. S. Sutton. "Learning to Predict by the Methods of Temporal Differences." *Machine Learning*, vol. 3, 9–44, 1988.
10. C. Watkins and P. Dayan. "Q-Learning." *Machine Learning*, vol. 8, no. 314, 279–292, 1992.
11. C. Boutilier. "Multiagent Systems: Challenges and Opportunities for Decision Theoretic Planning." *AI Mag.*, vol. 20, no. 35, 35–43, 1999.
12. C. Claus and C. Boutilier. "The Dynamics of Reinforcement Learning Cooperative Multiagent Systems. In *Proceedings of 15th National Conference on Artificial Intelligence*, pages 746–752, Madison, WI, June 1998.
13. J. Hu and M. P. Wellman. "Multiagent Reinforcement Learning: Theoretical Framework and an Algorithm." In *Proc. Fifteenth International Conference on Machine Learning*, pages 242–250, June 1998.
14. M. L. Littman. "Markov Games as a Framework for Multi-Agent Reinforcement Learning." In *Proceedings of the 11th International Conference on Machine Learning*, pages 157–163, 1994.

15. T. Sandholm and R. Crites. "Multiagent Reinforcement Learning in the Iterated Prisoner's Dilemma." *Biosystems*, vol. 37 147–166, 1995.
16. D. H. Wolpert and K. Tumer. An Introduction to Collective Intelligence. Technical Report NASA-ARC-IC-99-63, NASA Ames Research Center, 1999. URL:http://ic.arc.nasa.gov/ic/projects/coin_pubs.html. To appear in *Handbook of Agent Technology*, J. M. Bradshaw, (ed.) AAAI/MIT Press.
17. D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.
18. D. C. Parkes. *Iterative Combinatorial Auctions: Theory and Practice*. Ph.D. thesis, University of Pennsylvania, 2001.
19. N. Nisan and A. Ronen. "Algorithmic Mechanism Design." *Games Economic Behavior*, vol. 35, 166–196, 2001.
20. R. H. Crites and A. G. Barto. "Improving Elevator Performance Using Reinforcement Learning." In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems*—8, pages 1017–1023. MIT Press, 1996.
21. D. H. Wolpert. Theory of design of collectives. pre-print, 2002.
22. D. H. Wolpert. Bounded-rationality game theory. pre-print, 2001.
23. D. H. Wolpert, E. Bandari, and K. Tumer. Improving simulated annealing by recasting it as a non-cooperative game. 2001. submitted.
24. D. Monderer and L. S. Shapley. "Potential Games." *Games Economic Behavior*, vol. 14, 124–143, 1996.
25. D. H. Wolpert, K. Wheeler, and K. Tumer. "Collective Intelligence for Control of Distributed Dynamical Systems." *Europhysics Lett.*, vol. 49, no. 6, March 2000.
26. W. Nicholson. *Microeconomic Theory*. The Dryden Press, seventh edition, 1998.
27. D. H. Wolpert and K. Tumer. Collective Intelligence, Data Routing and Braess' paradox. *J. Artificial Intelligence Res.* 1995. to appear.
28. D. H. Wolpert, J. Sill, and K. Tumer. "Reinforcement Learning in Distributed Domains: Beyond Team Games." In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 819–824, Seattle, WA, 2001.
29. K. Tumer, A. Agogino, and D. Wolpert. "Learning Sequences of Actions in Collectives of Autonomous Agents." In *Proc. First International Joint Conference on Autonomous Agents Multi-Agent Systems*, Bologna, Italy, July 2002.
30. (a) M. P. Wellman. "A Market-Oriented Programming Environment and Its Application to Distributed Multicommodity Flow Problems." In *J. Art. Intel. Res.*, 1993.
30. (b) W. B. Arthur. "Complexity in Economic Theory: Inductive Reasoning and Bounded Rationality." *Amer. Econ. Rev.*, vol. 84, no. 2, 406–411, May 1994.
31. (a) D. Challet and Y. C. Zhang. "On the Minority Game: Analytical and Numerical Studies. *Physica A*, vol. 256 514, 1998.
31. (b) D. Wolpert and J. Lawson. "Designing Agent Collectives for Systems with Markovian Dynamics." In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy, July 2002.

BIBLIOGRAPHY

- C. Boutilier, Y. Shoham, and M. P. Wellman. "Editorial: Economic Principles of Multi-Agent Systems." *Artificial Intelligence J.*, vol. 94, 1–6, 1997.

- J. M. Bradshaw, (ed.), *Software Agents*. MIT Press, 1997.
- G. Caldarelli, M. Marsili, and Y. C. Zhang. "A Prototype Model of Stock Exchange." *Europhysics Let.*, vol. 40, 479–484, 1997.
- B. A. Huberman and T. Hogg. "The Behavior of Computational Ecologies." In *The Ecology of Computation*, pages 77–115. North-Holland, 1988.
- N. R. Jennings, K. Sycara, and M. Wooldridge. A Roadmap of Agent Research and Development. *Autonomous Agents Multi-Agent Sys.*, vol. 1, 7–38, 1998.
- N. F. Johnson, S. Jarvis, R. Jonson, P. Cheung, Y. R. Kwong, and P. M. Hui. "Volatility and Agent Adaptability in a Self-Organizing Market." preprint cond-mat/9802177, February 1998.
- T. Sandholm, K. Larson, M. Anderson, O. Shehory, and F. Tohme. "Anytime Coalition Structure Generation with Worst Case Guarantees." In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 46–53, 1998.
- S. Sen. *Multi-Agent Learning: Papers from the 1997 AAAI Workshop (Technical Report WS-97-03)*. AAAI Press, Menlo Park, CA, 1997.
- K. Sycara. "Multiagent Systems." *AI Magazine*, vol. 19, no. 2, 79–92, 1998.
- Y. C. Zhang. "Modeling Market Mechanism with Evolutionary Games." *Europhysics Let.*, March/April 1998.